

Express Mailing Label No.: ER211528536

PATENT APPLICATION

Docket No.: 1300.2.20

IBM Docket No.: RPS920030095US1

**UNITED STATES PATENT APPLICATION**

of

William W. Buchanan Jr.

Simon Chu

Linda Ann Riedle

and

Paul B. Tippet

for

**APPARATUS METHOD AND SYSTEM FOR  
ALTERNATE CONTROL OF A RAID ARRAY**

1                                   **APPARATUS METHOD AND SYSTEM FOR**  
2                                   **ALTERNATE CONTROL OF A RAID ARRAY**

3  
4                                   **BACKGROUND OF THE INVENTION**

5  
6           **1.     The Field of the Invention**

7           The invention relates to apparatuses, methods, and systems for controlling arrays of  
8           storage devices. Specifically, the invention relates to apparatus, methods, and systems for  
9           redundantly controlling a storage array.

10  
11          **2.     The Relevant Art**

12           RAID systems (*i.e.*, systems using Redundant Arrays of Independent Disks) are used  
13           to store large quantities of data within computer and storage networks. RAID systems are  
14           designed to be fault resistant and fault tolerant by distributing data among redundant arrays  
15           of independent disks usually with some form of error coding. Storage device redundancy and  
16           redundant data encoding facilitate continued operation and/or data recovery in the event of a  
17           failure within a particular storage device.

18           In addition to redundant storage devices, RAID systems often deploy redundant  
19           controllers to facilitate continued operation after a particular controller becomes inoperable.  
20           In addition to handling failures and controller redundancy properly, RAID controllers are  
21           typically required to receive write requests and associated data from a host computer,  
22           acknowledge reception of the data, stripe the data over a plurality of drives, and send the data  
23           to a storage array at a high throughput rate. As a result, RAID controllers are highly  
24           specialized and complex devices.

25           The specialized nature of RAID controllers, particularly redundant RAID controllers  
26           presents several challenges. In many redundant configurations, a secondary RAID controller  
27

1 waits in a standby or load sharing mode, ready to receive commands from the host computer  
2 in case a primary RAID controller is busy or fails. In some configurations, the secondary  
3 RAID controller performs no other function than waiting to take control in case of failure.  
4 The secondary RAID controller, although useful, adds significant cost to the RAID system.

5 RAID controllers are typically quite reliable in that they are typically solid state  
6 solutions with highly reliable components. In an attempt to reduce system cost, some  
7 administrators may forego using a redundant controller. Foregoing a controller is a risky  
8 proposition in that losses from downtime frequently exceed the cost of an entire system.  
9 Nevertheless, a demand for single controller solutions currently exists within data processing  
10 systems.

11 Furthermore, as specialized devices with specialized hardware to increase throughput,  
12 RAID controllers tend to fall out of date quickly – particularly as a system is updated and  
13 upgraded. Updating redundant RAID controllers requires purchasing the controllers in pairs.  
14 The added cost of the secondary RAID controller may hinder RAID system users from  
15 purchasing new controllers that could add significant functionality to their RAID systems. In  
16 addition, prolonging upgrades of old RAID controllers may place the controllers at additional  
17 risk for obsolescence or failure.

18 Another disadvantage of currently available storage control solutions is that, as solid  
19 state devices, RAID controllers are often fixed in their feature set at the time of manufacture  
20 and typically provide no facility for enhanced or augmented services such as advanced data  
21 recovery services, diagnostic services, firmware update services, and the like.

22 Given the disadvantages of currently available storage control solutions, what is  
23 needed is a cost-effective apparatus, method, and system for redundant control of a RAID  
24 system without using a dedicated secondary controller. Such an apparatus and method would  
25 reduce the cost of providing, upgrading, and augmenting RAID functionality within storage  
26 systems.

## SUMMARY OF THE INVENTION

The various elements of the present invention have been developed in response to the present state of the art, and in particular, in response to the problems and needs in the art that have not yet been fully solved by currently available redundantly controlled storage systems. Accordingly, the present invention provides an improved apparatus, method, and system for storage control of storage devices.

The present invention eliminates the need for a dedicated secondary controller by using a host-based adapter and associated control and emulation software executing on a host. In the context of the present specification, the term "dedicated controller" refers to a storage controller, particularly a RAID controller, with dedicated resources such as a CPU, logic unit, or the like. In particular, the present invention uses the shared resources of the host to reduce the cost and complexity of provided redundant control of a storage array or the like, while offering the additional advantage of augmented or enhanced services typically not found on currently available controllers.

In a first aspect of the present invention, an apparatus for storage control of a storage array includes a RAID controller operably connected to a storage array, a storage adapter residing on a host computer that is also operably connected to the storage array independent of the RAID controller, and a control module residing on the host computer. Due to the unique configuration of the present invention, the storage adapter is able to transmit storage commands to a selected storage device within the storage array.

In one embodiment, the control module selectively directs storage commands to the RAID controller and the storage adapter. During normal operation, the control module may direct commands supported by the RAID controller to the RAID controller and direct commands related to augmented or enhanced services to the storage adapter. In response to an inoperable RAID controller, the control module may also emulate the RAID controller

1 and conduct operations corresponding to storage commands on one or more selected storage  
2 devices within the storage array.

3 As described in the present invention the storage adapter is an electrical interface  
4 between the control module and the storage devices. Possible implementations of the storage  
5 adapter include a SCSI adapter, an SSA adapter, a fiberchannel adapter, an IDE adapter, or  
6 the like. Using a standard storage adapter such as the aforementioned adapters often  
7 facilitates using the present invention without purchasing additional hardware.

8 A controller storage adapter may also provide an interface between the control  
9 module and the RAID controller. However, in certain embodiments, the RAID controller  
10 does not interface with a storage adapter, but rather resides on the host computer. By  
11 including the RAID controller on the host computer, the present invention may provide  
12 improved performance in a more cost-effective manner than previous solutions.

13 There are many possible different implementations for the RAID controller. Some  
14 implementations include a SCSI RAID controller, an SSA RAID controller, a fiberchannel  
15 RAID controller, an IDE RAID controller, or the like. Those skilled in the art may also  
16 derive other implementations for the RAID controller.

17 In certain embodiments, the control module selectively emulates a RAID controller.  
18 By emulating a RAID controller, the control module may conduct RAID operations as well  
19 as non-RAID operations on the storage devices via the storage adapter. In one embodiment,  
20 the supported operations include read operations, write operations, recovery operations,  
21 diagnostic operations, formatting operations, and firmware update operations. In one  
22 embodiment, the control module is a device driver that translates operating system calls to  
23 storage device commands.

24 In one embodiment, the control module also includes a multipath software layer that  
25 selects a control path through either the RAID controller or the storage adapter. The  
26 multipath software layer may also select multiple control paths in parallel to allow the control  
27

1 module to perform augmented services on the storage devices while the RAID controller is  
2 performing standard services on the storage devices. Thus, the present invention provides  
3 cost-effective enhanced functionality to the RAID system.

4 In another aspect of the present invention, a method for controlling a storage array  
5 includes transmitting storage commands to a RAID controller that conducts operations on a  
6 storage array, transmitting storage commands via a storage adapter to a selected storage  
7 device of the storage array, and selectively directing storage commands to the RAID  
8 controller and the storage adapter.

9 In one embodiment, the method further includes emulating a RAID controller. By  
10 emulating the RAID controller, the method may provide alternate control of the storage array  
11 to a host computer or the like. In addition, enhanced or augmented services as well as  
12 backup or secondary control of the storage devices may be provided by the aforementioned  
13 method.

14 Various elements of the present invention are combined into a storage system of the  
15 present invention including a host computer, a storage array, a RAID controller operably  
16 connected to the storage array, a storage adapter operably connected to the storage array  
17 independent of a storage controller, and a control module residing on the host computer.

18 The control module of the present invention selectively directs storage commands to  
19 the RAID controller and the storage adapter. In response to receiving storage commands  
20 from the control module, the RAID controller executes corresponding operations on the  
21 storage array. Additionally, the control module may transmit storage commands from the  
22 control module to a selected storage device within the storage array via the storage adapter.

23 In another aspect of the present invention, a method for servicing a storage array  
24 includes operably connecting a storage array to a storage adapter independent of a RAID  
25 controller and installing a control module on the host computer. The method for servicing a  
26  
27

1 storage array enables a technician or the like to provide diagnostic operations, data recovery,  
2 firmware updates and other services related to storage arrays.

3 The method for servicing a storage array may also include operably connecting a  
4 storage adapter to a storage array, and transmitting storage commands to one or more  
5 selected storage devices within the storage array. The method may also include disabling or  
6 quiescing the RAID controller and transferring control of the storage array to the storage  
7 adapter. Transferring control to the storage adapter allows the method to use an alternate  
8 control path to communicate with the storage devices. In communicating with the storage  
9 devices, the method may update firmware within a storage device, conduct diagnostic  
10 operations, conduct recovery operations or the like.

11 The various elements and aspects of the present invention provide a cost-effective  
12 alternative to a dedicated redundant RAID controller. Additionally, the present invention  
13 provides augmented services and features not found in currently available RAID storage  
14 systems. These and other features and advantages of the present invention will become more  
15 fully apparent from the following description and appended claims, or may be learned by the  
16 practice of the invention as set forth hereinafter.

## **BRIEF DESCRIPTION OF THE DRAWINGS**

In order that the manner in which the advantages and objects of the invention are obtained will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof, which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

Figure 1 is a schematic block diagram illustrating a typical prior art storage system;

Figure 2 is a schematic block diagram illustrating one embodiment of an alternate control storage system of the present invention;

Figure 3 is a schematic block diagram illustrating an additional embodiment of the alternate control storage system of the present invention;

Figure 4 is a schematic block diagram illustrating one embodiment of a storage control apparatus of the present invention;

Figure 5 is a flow chart diagram illustrating one embodiment of a storage control method of the present invention;

Figure 6 is a flow chart diagram illustrating a storage system service method of the present invention; and

Figure 7 is a schematic block diagram illustrating selected embodiments of an interconnection network in accordance with the present invention.



## DETAILED DESCRIPTION OF THE INVENTION

Many of the functional units described in this specification have been labeled as modules, in order to more particularly emphasize their implementation independence. For example, a module may be implemented as a hardware circuit comprising custom VLSI circuits or gate arrays, off-the-shelf semiconductors such as logic chips, transistors, or other discrete components. A module may also be implemented in programmable hardware devices such as field programmable gate arrays, programmable array logic, programmable logic devices or the like.

Modules may also be implemented in software for execution by various types of processors. An identified module of executable code may, for instance, comprise one or more physical or logical blocks of computer instructions which may, for instance, be organized as an object, procedure, or function. Nevertheless, the executables of an identified module need not be physically located together, but may comprise disparate instructions stored in different locations which, when joined logically together, comprise the module and achieve the stated purpose for the module.

Indeed, a module of executable code could be a single instruction, or many instructions, and may even be distributed over several different code segments, among different programs, and across several memory devices. Similarly, operational data may be identified and illustrated herein within modules, and may be embodied in any suitable form and organized within any suitable type of data structure. The operational data may be collected as a single data set, or may be distributed over different locations including over different storage devices, and may exist, at least partially, merely as electronic signals on a system or network.

Figure 1 is a schematic block diagram illustrating a typical prior art storage system 100. The depicted storage system 100 includes a host 110, an adapter 114, a redundant set of RAID controllers 120, an interconnection network 132, and an array of storage devices 134.

1 In the depicted storage system 100, the host 110 sends commands to the RAID controllers  
2 120. In turn the RAID controllers 120 may communicate with the storage devices 134 via  
3 the interconnection network 132.

4 In one redundant arrangement, a primary RAID controller 120a executes operations  
5 on the storage devices 134 in response to commands issued by the host 110. In the  
6 aforementioned arrangement, a secondary RAID controller 120b may be activated in  
7 response to a failure by the primary RAID controller 120a. By activating the secondary  
8 RAID controller 120b in response to failure of the primary RAID controller 120a, reliability  
9 and availability of the storage system 100 is improved.

10 The cost of achieving the improved reliability of the storage system 100 may be  
11 substantial. The RAID controllers 120 are typically non-commodity specialized devices that  
12 are expensive to purchase, install, and maintain. For example, the RAID controllers often  
13 contain components that are highly specialized, quickly out-dated, and more expensive than  
14 general purpose storage controllers.

15 Figures 2 and 3 are schematic block diagrams illustrating selected embodiments of an  
16 alternate control storage system 200 of the present invention. As depicted, the alternate  
17 control storage system 200 includes a host 110, a primary RAID controller 120a, an  
18 interconnection network 132, one or more storage devices 134, a control module 212, one or  
19 more adapters 214, an internal bus 216, and multiple interface ports 218. The alternate  
20 control storage system 200 is a low-cost, resource-efficient alternative to currently available  
21 redundant storage systems.

22 Figure 2 depicts a particular embodiment of the alternate control storage system 200,  
23 namely an alternate control storage system 200a, wherein the primary RAID controller 120a  
24 is external to the host 110. Likewise, Figure 3 depicts a particular embodiment of the  
25 alternate control storage system 200, namely an alternate control storage system 200b,  
26 wherein the primary RAID controller 120a resides on the host 110.

1 The host 110 may interface to the primary RAID controller 120a via the internal bus  
2 216 as depicted in Figure 2, or via one of the interface ports 218 as depicted in Figure 3.

3 In turn, the primary RAID controller 120a communicates with the storage devices 134 via the  
4 interconnection network 132. To facilitate communication with the interconnection network  
5 132, the primary RAID controller 120a may include one or more interface ports 218 that  
6 provide a storage interface or storage network interface such as a SCSI interface, an SSA  
7 interface, a fiberchannel interface, an IDE interface, an ethernet interface, or the like.

8 In the depicted alternate control storage system 200a, the control module 212 and the  
9 adapters 214 reside on the host 110. One of the adapters 214 may be a controller adapter  
10 214a while another may be a storage adapter 214b. During normal operation, the control  
11 module 212 directs standard storage commands to the primary RAID controller 120a via the  
12 controller adapter 214a. In turn, the primary RAID controller 120a executes operations on  
13 the storage devices 134.

14 In one embodiment, in response to failure of the primary RAID controller 120a, the  
15 control module 212 emulates the primary RAID controller 120a by receiving commands  
16 from the host and executing corresponding operations on the storage devices 134 via the  
17 storage adapter 214b and the interconnection network 132. Thus, the control module 212  
18 and the storage adapter 214b essentially emulate a RAID controller 120, eliminating the need  
19 to include an expensive secondary RAID controller 120b.

20 In one embodiment, the adapters 214, including the controller adapter 214a and the  
21 storage adapter 214b, may be SCSI adapters, SSA adapters, fiberchannel adapters, IDE  
22 adapters, or the like. The aforementioned adapters 214 are typically significantly less  
23 expensive than the RAID controllers 120.

24 In one embodiment, the control module 212 disables or quiesces the primary RAID  
25 controller 120a in response to failure of the primary RAID controller 120a. The control  
26 module 212 may also include a driver configured to translate operating system calls into  
27

1 storage commands. These operations may include augmented or enhanced services such as  
2 recovery operations, diagnostic operations, firmware update operations, and the like.

3 In the depicted alternate control storage system 200b, the primary RAID controller  
4 120a resides on the host 110 and receives storage commands via the internal bus 216. In one  
5 embodiment, the primary RAID controller 120a is a peripheral card such as a PCI card, that  
6 is inserted into a slot to electrically connect to the internal bus 216. In another embodiment,  
7 the primary RAID controller 120a is integrated onto a motherboard residing on the host 110  
8 and directly connected to the internal bus 216. In either arrangement, placing the primary  
9 RAID controller 120a within the host 110 may eliminate interface hardware associated with  
10 the external primary RAID controller 120a, increase system performance, and reduce the cost  
11 of providing redundancy within a storage system.

12 Figure 4 is a schematic block diagram illustrating one embodiment of a storage  
13 control apparatus 400 of the present invention. As depicted, the storage control apparatus  
14 400 includes a control module 212, a primary RAID controller 120a, a storage adapter 214b,  
15 a path selection module 410, a RAID controller emulation module 412, a failure detection  
16 module 414, and an augmented services module 416. The control module 212 facilitates  
17 redundant control of a storage array using a single RAID controller 120a.

18 The path selection module 410 directs storage commands to either the primary RAID  
19 controller 120a or the storage adapter 214b. In one embodiment, the path selection module  
20 410 is essentially a multipath software layer within the system code of the host. During  
21 normal operation, the path selection module 410 directs standard storage commands received  
22 from the host 110 to the primary RAID controller 120a.

23 In one embodiment, the failure detection module 414 informs the path selection  
24 module 410 when the primary RAID controller 120a is not functioning properly. In response  
25 to the failure, the path selection module 410 directs commands received from the host 110 to  
26 the RAID controller emulation module 412. When activated, the emulation module 412 may  
27

1 translate high-level storage commands including RAID commands into corresponding low-  
2 level storage commands supported by the storage devices 134.

3 During normal operation, the path selection module 410 may also direct requests for  
4 augmented or enhanced services to the augmented services module. The augmented services  
5 module 416 performs services on the storage devices 134 that are not available on the RAID  
6 controller 120a.

7 In one embodiment, the primary RAID controller 120a and the storage adapter 214b  
8 may both communicate with the storage devices 134 simultaneously through managed  
9 synchronization. While the primary RAID controller 120a is executing operations on the  
10 storage devices 134, the storage adapter 214b may also perform augmented services on the  
11 storage devices 134. In other embodiments, the RAID controller 120a is put on standby or is  
12 disabled during service operations conducted by the storage adapter 214b.

13 Figure 5a is a flow chart illustrating one embodiment of a failure response method  
14 500 of the present invention. The storage control method 500 may be conducted in  
15 conjunction with the storage control apparatus 400, the storage system 200a, 200b, or the  
16 like. As depicted, the failure response method 500 includes a failure detected test 510, a  
17 retrieve configuration step 520, an initialize emulator step 530, and a shutdown RAID  
18 controller step 540. The failure response method 500 facilitates fault tolerant handling of  
19 RAID control failures.

20 While depicted in a certain sequential order, the depicted functionality of the failure  
21 response method 500 may be conducted in various sequences according to the needs and  
22 capabilities of the particular embodiment. For example, in one embodiment the depicted  
23 functionality is invoked by system events such that the order of execution is event dependent.

24 The failure detected test 510 ascertains whether a failure has occurred to a RAID  
25 controller. In one embodiment, status query messages are sent to the RAID controller and a  
26 timeout event is set in case the RAID controller does not respond. In another embodiment,

1 status registers on the RAID controller are accessed to assess the status of the controller. In  
2 yet another embodiment, a test sequence such as sending a write command to the RAID  
3 controller and accessing striping information on individual storage devices is conducted in  
4 order to ascertain if the RAID controller is functioning properly.

5 The retrieve configuration step 520 retrieves RAID configuration information such as  
6 RAID level and striping information from local configuration files, system configuration  
7 files, registers within one or more RAID controllers, or similar sources. The RAID  
8 configuration information may facilitate providing augmented services or alternate control of  
9 a RAID array. For example, the RAID configuration information may facilitate monitoring a  
10 RAID controller, detecting RAID controller failures and emulating a RAID controller.

11 The initialize emulator step 530 initializes a controller emulator such as a RAID  
12 controller emulator. In one embodiment, the initialize emulator step 530 uses the  
13 configuration information retrieved at step 510 to initialize the controller emulation module  
14 depicted in Figure 4.

15 The shutdown controller step 540 attempts to shutdown a faulty RAID controller, or  
16 the like. In one embodiment, the controller is quiesced to clear the processing queues within  
17 the controller. Subsequently data caches on the controller may be flushed to clear any write  
18 data within data caches, and followed by invocation of a disable command to prevent further  
19 processing by the controller. In response to completion of the shutdown controller step 540  
20 the depicted failure response method 500 ends 545.

21 Figure 5b is a flow chart illustrating one embodiment of a storage control method 550  
22 of the present invention. The storage control method 550 may be conducted in conjunction  
23 with the storage control apparatus 400, the storage system 200a, 200b, or the like. The  
24 depicted embodiment of the storage control method 550 includes a receive command step  
25 555, an augmented service test 560, a perform augmented service step 565, a controller  
26 functional test 570, a perform standard service step 575, and a transmit command step 580.

1 The method begins with the receive command step 555. In one embodiment, a  
2 storage command or the like is received from a storage manager or the like executing on the  
3 host 110. Upon reception of the storage command, the method proceeds to the augmented  
4 service test 560. At the augmented service test 560, the method 500 determines whether the  
5 command corresponds to an augmented or enhanced service.

6 If the received command corresponds to an augmented or enhanced service, the  
7 method proceeds to the perform augmented service step 565. In one embodiment, the  
8 augmented services include disabling or quiescing the RAID controller, conducting  
9 diagnostic operations, recovery operations, firmware updates, or the like. Upon completion  
10 of the perform augmented service step 565, the storage control method 550 loops to the  
11 receive command step 555.

12 If an augmented or enhanced service was not requested, the method continues with  
13 the controller functional test 570. At the controller functional test, 570 the method  
14 determines whether a dedicated controller (such as a RAID controller) is functional. If the  
15 dedicated controller is not functional, the method proceeds to the perform standard service  
16 step 575. In one embodiment, standard services are performed by an emulation module and  
17 include conducting RAID operations on a storage array. Upon completion of the perform  
18 augmented service step 565, the storage control method 550 loops to the receive command  
19 step 555.

20 If the controller is functional, the storage control method 500 proceeds to the transmit  
21 command step 580 in order to invoke the supported service within the RAID controller or the  
22 like. Thereafter, the method 550 returns to the receive storage command step 555. The  
23 depicted storage control method 550 facilitates providing alternate control of a storage array  
24 in the event of a failure of a dedicated controller while also providing host-based augmented  
25 storage services not available on the dedicated controller.

Figure 6 is a flow chart illustrating one embodiment of a storage system service method 600 of the present invention. The depicted storage system service method 600 includes an install service module step 610, a connect storage adapter step 620, a transfer control step 630, a service controller step 640, a recover data step 650, and a service storage devices step 660.

The depicted storage system service method 600 may be performed by a service technician or the like in order to install or service the alternate control system 200 or the like.

While depicted in a certain sequential order, the depicted steps may be performed in a manner or order appropriate to the particular system being serviced. For example, the service controller step 640 and the recover data step 650 may be omitted in systems wherein the steps are not needed.

The install service module step 610 installs a service module, or the like, on the host 110. In one embodiment, the service module 610 is essentially the control module 212 with additional functionality appropriate to installing or servicing the alternate control storage system 200, or the like.

The installed service module may perform or direct completion of the storage system service method 600. At the connect storage adapter step 620, a storage adapter such as the storage adapter 214b is connected to a storage array such as the storage devices 134 via an interconnection network such as the interconnection network 132. In one embodiment, the service module (not shown) executing on the host computer indicates which storage devices 134 are accessible via the storage adapter 214b. The storage adapter 214b may be an adapter newly installed in conjunction with step 620 or an adapter previously available on the host 110.

In conjunction with the install service module step 610 and the connect storage adapter step 620, RAID configuration information such as RAID level and striping information may be retrieved from system configuration files, registers within one or more



1 RAID controllers, or similar sources. The RAID configuration information may facilitate  
2 providing augmented services or alternate control of a RAID array. For example, the RAID  
3 configuration information may facilitate monitoring a RAID controller, detecting RAID  
4 controller failures and emulating a RAID controller.

5 The transfer control step 630 transfers control of the storage devices to the service  
6 module (not shown), the control module 212, or the like. In one embodiment, transferring  
7 control includes quiescing or disabling a dedicated controller such as the RAID controller  
8 120a. Transferring control may facilitate continued operation while performing diagnostic or  
9 maintenance services on the storage system.

10 The service controller step 640 services a dedicated controller such as the RAID  
11 controller 120a. In one embodiment, servicing the dedicated controller includes updating  
12 firmware on the controller. In another embodiment, servicing the dedicated controller  
13 includes replacing the controller with a different controller.

14 The recover data step 650 performs data recovery operations. In one embodiment, the  
15 data recovery operations restore images on the storage devices 134 to a desired state. In  
16 another embodiment, the data recovery operations transfer data within obsolete storage  
17 devices to newly installed storage devices.

18 The service storage devices step 660 services the storage devices 134 or the like. In  
19 one embodiment, servicing the storage devices 134 includes updating firmware associated  
20 with the storage devices 134. In another embodiment, servicing the storage devices 134  
21 includes conducting diagnostic operations.

22 Figure 7 is a schematic block diagram illustrating selected embodiments of the  
23 interconnection network 132 in accordance with the present invention. The depicted  
24 embodiments of the interconnection network 132 are intended to illustrate without limitation  
25 various interconnection technologies and topologies that may be used with the present  
26 invention. The interconnection network 132 may include a variety of elements appropriate to  
27

1 interconnecting storage devices such as point-to-point connections 710, storage device  
2 connections 720, controller connections 730, redundant controller connections 740, parallel  
3 buses 750, hubs 760, and hub interconnections 770.

4 The controller connection 730 provides operable connectivity such as electrical,  
5 optical, or wireless connectivity from and to a dedicated controller, such as the RAID  
6 controller 120a. The redundant connection 740 provides operable connectivity such as  
7 electrical, optical, or wireless connectivity from and to a host-based storage adapter such as  
8 the storage adapter 214b. The storage device connections 720 provide operable connectivity  
9 such as electrical, optical, or wireless connectivity from and to storage devices such as the  
10 storage devices 134 depicted in various Figures.

11 The selected embodiments shown in Figure 7 include a point-to-point interconnection  
12 network 132a, a bused interconnection network 132b, and a storage area network 132c. The  
13 depicted point-to-point interconnection network 132a includes a set of point-to-point  
14 connections 710, such as fiberchannel connections or the like, that operably connect the  
15 controller connection 730, the redundant controller connection 740, and the storage device  
16 connections 720. In one embodiment, the point-to-point connections 710 form a loop of bi-  
17 directional connections such that the interconnection network 132a remains functional in the  
18 event of a device or connection failure.

19 The depicted bused interconnection network 132b includes a parallel bus 750. The  
20 parallel bus 750 operably connects the controller connection 730, the redundant controller  
21 connection 740, and the storage device connections 720. The parallel bus 750 enables  
22 operative commands to be sent to each storage device in parallel. Through the controller  
23 connection 730, the dedicated controller, such as a RAID controller, sends operative  
24 commands. The operative commands travel through the parallel bus 750 to the storage  
25 device connections 720. Storage devices such as the storage devices 134 receive the  
26 operative commands from the storage device connections 720.

1 The depicted storage area network 132c includes one or more hubs 760 connected by  
2 one or more hub interconnects 770. In the depicted embodiment, two hubs 760 are shown  
3 connected by one hub interconnect 770. The hub 760a is operably connected to the  
4 controller connection 730 and the storage device connections 720. The hub 760b is operably  
5 connected to the redundant controller connection 740 and additional storage device  
6 connections 720. The hub 760a receives operative commands from the dedicated controller,  
7 such as a RAID controller, from the controller connection 730. The operative commands  
8 pass to the storage device connections 720 via hub 760a and the hub 760b.

9 The present invention provides secondary or augmented control of a storage array in a  
10 cost-effective manner. The present invention may be embodied in other specific forms  
11 without departing from its spirit or essential characteristics. The described embodiments are  
12 to be considered in all respects only as illustrative and not restrictive. The scope of the  
13 invention is, therefore, indicated by the appended claims rather than by the foregoing  
14 description. All changes which come within the meaning and range of equivalency of the  
15 claims are to be embraced within their scope.

16 What is claimed is:  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27